

UNITED STATES PATENT APPLICATION

For

COMMUNICATION SYSTEM AND METHOD FOR MEDIA ACCESS CONTROL

Inventors:

Erez Bashan, Atlanta, GA.
Erez Izenberg, Tel Aviv, IL

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP
12400 Wilshire Boulevard
Los Angeles, CA 90025-1026
(408) 947-8200

Attorney's Docket No.: 6985P001

"Express Mail" mailing label number: EV410137653US

Date of Deposit: February 20, 2004

I hereby certify that I am causing this paper or fee to be deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and that this paper or fee has been addressed to the Commissioner for Patents, P.O. Box 1450 Alexandria, VA 22313-1450

Jennifer L. Stewart

(Typed or printed name of person mailing paper or fee)

Jennifer L. Stewart

(Signature of person mailing paper or fee)

February 20, 2004

(Date signed)

COMMUNICATION SYSTEM AND METHOD FOR MEDIA ACCESS CONTROL

RELATED APPLICATION

[0001] The present application is related to, claims the priority benefit of and incorporates by reference Israeli Patent Application 154560, filed February 20, 2003.

FIELD OF THE INVENTION

[0002] The present invention relates to communication systems and methods in general, and to methods and systems for media access control of a shared communication media. Even more particularly, to methods and systems for controlling upstream transmission of a point to multipoint optical network. The invention is applicable to, but not limited to, point to multipoint passive optical networks.

BACKGROUND OF THE INVENTION

[0003] Optical access networks, such as point to multipoint optical access networks are known in the art. ITU-T Recommendation G.983.1 defines an access network that utilizes optical fiber technology to convey point to multipoint downstream traffic from an headend such as an Optical Line Termination (OLT) to multiple network units such as Optical Network Units (ONUs) or Optical Network Terminations (ONTs) and to convey point to point upstream traffic from an ONU or ONT to the OLT.

[0004] An ONU provides, either directly or remotely, the user-side interface of the Optical Access Network (OAN).

[0005] An OLT provides the network side interface of the OAN and is connected to at least one Optical Distribution Network (ODN).

[0006] An ODN provides the optical transmission means between the OLT and the ONUs. An ODN usually includes passive optical components such as optical fibers, optical connectors, passive branching components, passive optical attenuators and splices.

[0007] The optical access network supports Asynchronous Transfer Mode (ATM) based data transmission. ATM based data transmission allows to support more than a single class of service. When a packet, such as but not limited to Internet Protocol (IP) packet, is received at an ATM supporting network, it is segmented to provide a group of at least one ATM cell. Each ATM cell is routed across the ATM based network. Before exiting the ATM based network the IP packet is reassembled from the at least one ATM cell.

[0008] Each ATM cell is 53 bytes long and includes a 7-byte header and 48-byte payload.

The ATM header is utilized for routing ATM cells across the ATM based network.

[0009] As multiple network units, such as ONUs or ONTs share the same media, the OLT controls upstream transmission, by assigning upstream timeslots by implementing a Media Access Control (MAC) scheme. According to the ITU-T Recommendation G.983.4 upstream bandwidth may assigned in two manners – (a) in response to the utilization of upstream bandwidth by each of the ONUs, and (b) in response to upstream status reports from the ONUs or ONTs. More specifically, each ONU (or ONT) can include at least one Transmission Container (T-CONT), each T-CONT has at least one queue. An ONU reports the queue length of T-CONTs that belong to him. Usually, a single T-CONT has more than a single queue, each queue associated with a distinct class of service. Accordingly, the aggregate queues length of that T-CONT is reported.

[00010] An ONU (or ONT) can selectively report the T-CONTs queue length by various manners. For example, An ONU (or ONT) can report the queues lengths from the most congested T-CONT, from each T-CONT in turn equally, but this is not necessarily so. These reports are transmitted upstream in mini-slots assigned by the OLT.

[00011] Upstream traffic is arranged in an upstream frame of 53 timeslots. Each timeslot consists of three-bytes of PON layer overhead and either an ATM cell or a PLOAM cell.

[00012] The OLT allocates upstream bandwidth on a timeslot to timeslot basis, according to the T-CONTs queue length and then transmits downstream data grants in downstream PLOAM cells. Assuming that the upstream and downstream bit rate are the same, then during a frame of 53 timeslots, two PLOAM cells are utilized for providing 53 data grants, corresponding to the 53 timeslots within each upstream frame. When the upstream data rate is much smaller than the downstream data rate, some PLOAM cells are empty.

[00013] The T-CONTs are classified to five types, each type is characterized by an assigned bandwidth out of the following five assigned bandwidth types: (i) fixed bandwidth, (ii) assured bandwidth, (iii) non-assured bandwidth, (iv) best effort bandwidth, and (v) maximum bandwidth. It is noted that the first four assigned bandwidths are listed according to their priority, starting with the highest priority assigned bandwidth. Accordingly, the assignment of bandwidth starts by assigning bandwidth to fixed bandwidth. The assignment limits the amount of cell transfer delay and delay variation. Assured bandwidth is assigned using the remaining bandwidth. Assured bandwidth means that a predefined average (long- range) bandwidth is assigned. It is noted that

the amount of allocated bandwidth per frame can fluctuate. The yet remaining bandwidth (also referred to as surplus bandwidth) is utilized for the lower priority bandwidth assignments such as the non-assured bandwidth and the best effort bandwidth.

[00014] A type 1 T-CONT is characterized by fixed bandwidth only. Bandwidth is allocated to a type 1 T-CONT regardless of whether its queues are empty or not. A type 2 T-CONT is characterized by assured bandwidth only.

[00015] A type 3 T-CONT is characterized by assured bandwidth and non-assured bandwidth. A type 3 T-CONT shall be allocated bandwidth equivalent to its Assured bandwidth, only when it has cells at a rate equivalent to Assured bandwidth or more than Assured bandwidth. Non-assured bandwidth shall be allocated across T-CONTs with Assured bandwidth, and by requesting surplus bandwidth in proportion to the Assured bandwidth of the individual T-CONT on the PON, e.g., Weighted Round Robin method, as surplus bandwidth. The sum of the assured bandwidth and non-assured bandwidth allocated to this T-CONT should not exceed its maximum bandwidth, which is pre-provisioned.

[00016] A type 4 T-CONT is characterized by best effort bandwidth only and does not have any guaranteed bandwidth. A type 4 T-CONT shall only use bandwidth that has not been allocated as fixed bandwidth, assured bandwidth and non-assured bandwidth to other T-CONTs sharing the same upstream bandwidth. Best-effort bandwidth is allocated to each type 4 T-CONT equivalently, e.g., based on Round Robin method, up to their predefined maximum bandwidth.

[00017] A type 5 T-CONT is the super set of type 1 – type 4 T-CONTs. Accordingly they can be characterized by at least one of the following assigned bandwidths: fixed bandwidth, assured bandwidth, non-assured bandwidth and best-effort bandwidth. It is noted that the bandwidth allocation cannot exceed the maximum bandwidth of the T-CONT.

[00018] It is noted that a T-CONT may have a priority control mechanism and/or an internal schedule that are operable to determine from which class of service queue to transmit an ATM cell in response to a data grant.

[00019] It is further noted that each assigned bandwidth type may be associated with its own class of service. Accordingly, the five mentioned above type of bandwidth assignment are associated with five class of service.

[00020] U.S. patent number 5,926,478 of Ghaibeh et al. titled “Data transmission over a point-to-multipoint optical network” describes a data transmission protocol for use in an ATM-based point-to-multipoint passive optical network interconnecting a headend facility and a

plurality of network units. The headend facility controls the upstream transmission from the network units in response to ATM cell queue sizes at the network units and in response to a selected set of service priorities. A network unit can include various queues, such as a CBR queue, a VBR queue and a ABR queue. The sizes of these queues are included within an upstream bandwidth request.

[00021] Downstream data is transmitted in serial data frames comprising one hundred eighty, fifty-four byte downstream slots, including two framing slots and one hundred seventy-eight ATM cell slots. Each downstream frame slot includes a one byte MAC overhead header field for transmitting upstream transmission permits allocated over twenty bit permit fields, for a total of seventy-two upstream permits allocated per downstream frame. The downstream frames are transmitted every 125 .mu.sec for an overall downstream bit rate of 622.08 Mbps. Upstream data is transmitted from an individual network units in five hundred forty bit upstream data slots, each upstream slot having a preamble portion and a payload portion, i.e., with seventy-two upstream slots are transmitted every 125 .mu.sec, thereby forming upstream frames received at the headend at a data rate of 311.04 Mbps.

[00022] End users negotiate with service providers to determine a class of service or service level. A service Level Agreement (SLA) defines traffic parameters from the end-user, throughout at least one network that interconnect the end-user with other end-users or with other service providers. The traffic parameters include overall delay, delay fluctuations, bandwidth allocation and the like. Many users generate and receive variable length packets, such as Internet Protocol (IP) packets. In such cases the SLA relates to the transmission of IP packets, and does not necessarily comply with the transmission and routing of ATM cells originating from the IP packets.

[00023] The size of the IP packet is smaller then the aggregate size of ATM cells the originate from the IP packet, as a header of 7 bytes is added to each 48 bytes of IP packet. Assuming that the size of the IP network is $S1 = A * 48 + B$, then the aggregate size of the ATM cells that originate from the IP packet ($S2$) equals: $S2 = 53 * (\text{trunc}\{S1/ 48 \text{ byte}\} + 1) = 53*(A+1)$. It is noted that the $(A+1)$ 'th ATM cell includes B bytes originating from the IP packet and $(48 - B)$ "stuffing" bytes. The utilization of the ATM network equals $S2/S1$. $S1$ is also referred to as the length of the relevant payload and $(S2-S1)$ is also referred to as the overhead signals. ATM network arbitration and scheduling schemes are based upon the overall aggregate size of ATM cells (including header and stuffing bit) that are stored within queues, and not according to the aggregate "net" payload of the cells and neither upon the length of each group.

[00024] Furthermore, ATM policing schemes, such as RED or WRED schemes that determine when to discard incoming traffic in response to various parameter such as traffic load are also cell oriented. Accordingly, most of the ATM cells originating from the same IP packet can be routed across an ATM network just to see that one ATM cell was discarded, thus all the ATM cells must be re-routed across the network.

[00025] As ATM cells are routed on a cell to cell basis, in the presence of ATM cells from many sources, consecutive cells originating from the same IP packet are not routed in a consecutive manner, thus increasing the overall delay and the delay jitter across the ATM network, and require extensive allocation of memory resources.

[00026] There is a need for a system and method for improved bandwidth utilization.

[00027] There is a need for a system and method for a media access control method and controller that are based upon relevant payload.

[00028] There is a need for a system and method and for allowing the routing of ATM cell groups.

SUMMARY OF THE PRESENT INVENTION

[00029] The invention provides a communication system that includes an optical communication network, interconnecting a headend and a plurality of network units; wherein the headend has a media access controller for issuing data grants and grouping information requests; wherein a data grant being issued at least partially in response to previously received grouping information. At least some network units out of the plurality of network units are operable to: (i) receive data to be upstream transmitted to the headend (ii); upstream transmit grouping information associated with the received data (iii); and upstream transmit data to the headend in response to data grants issued by the media access controller.

[00030] A data grant authorizes an identified network unit out of the plurality of network units to upstream transmit a group of consecutive data cells during at least one consecutive timeslot.

[00031] According to an aspect of the invention the optical communication network is a passive optical network, the headend is an OLT and the network units are ONTs and/or ONUs but can also be T-CONTs. A single ONT or a single ONU can include a variety of T-CONTs. Conveniently, the communication network is ITU-T Recommendation G.983.4 compliant, and includes a segmentation and reassembly units for converting IP packets to groups of ATM cells. Preferably, a group of ATM cells originate from a single IP packet, and the grouping information reflects the length of the IP packet.

[00032] The media access controller can include a plurality of arbitrators. Each arbitrator may be associated with a class of service out of a plurality of class of services. The class of services may include: provisioned bandwidth, guaranteed bandwidth, assured bandwidth and non-assured bandwidth, but may also include at least one of the following class of services: provisioned bandwidth, minimum latency, assured bandwidth, non-assured bandwidth, minimum drop and minimum jitter. An arbitrator that is associated with a certain class of service is operable to arbitrate between transmission requests from queues of the same class of service, but this is not necessarily so, as a single arbitrator may also be operable to arbitrate between queues of distinct class of service.

[00033] The headend is operable to transmit data to network units in consecutive frames; wherein each frame further includes at least one data grant. Each frame may further include a grouping information request. Usually, each frame includes a plurality of fixed length slots.

[00034] The invention provides a system wherein grouping information includes grouping information units. A grouping information unit may represent a parameter of a group of data cells that are stored in a queue within a network unit. Conveniently, each group of data cells may include relevant payload and overhead signals. Preferably, a grouping information unit reflects the length of the relevant payload. It is further noted that the queue may store a packet, whereas the segmentation of a packet to at least one cell occurs only after the upstream transmission is approved.

[00035] The invention provides a system wherein the media access controller is operable to determine an amount of grouping information to be sent from a network unit. The determination may be responsive to grouping information previously transmitted from the network unit and to a data threshold, but can also be responsive to an estimation of grouping information relating to information that is yet to be sent upstream from the network units. The data threshold conveniently reflects a maximal amount of data that can be upstream transmitted from the network unit to the headend during a predefined time period. Conveniently, the predefined time period equals a MAC cycle, of an upstream frame period.

[00036] The invention provides a system wherein at least some of the network units are not operable to generate grouping information. In such a case the media access controller estimates grouping information relating to data from said network units.

[00037] The invention provides a system wherein at least some network units out of the plurality of network units include (i) a first input port for receiving variable length data packets and (ii) a segmenting and grouping information unit that is operable to (ii.a) segment a received variable length data packet to provide a group of fixed sized data cells, and to (ii.b) generate grouping information reflecting a parameter of the group of fixed sized cells. Conveniently, the at least some network units further include a classifier, for classifying incoming data packets in response to their service of class. Preferably, the variable length data packets are Internet Protocol packets and the fixed sized cells are Asynchronous Transfer Mode cells.

[00038] The invention provides a system wherein at least some network units out of the plurality of network units include (i) a second input port for receiving fixed sized cells, and (ii) an assembly unit for grouping the fixed sized cells to fixed sized cell groups. Conveniently, the at least some network unit further include at least one of the following units: (iii) a grouping information generator, for generating grouping information representative of a parameter of a group of fixed sized cells, or (iv) a grouping information extractor, for extracting grouping information embedded within the group of fixed sized cells.

[00039] The invention provides a system wherein the media access controller is operable to provide data grants in response to at least one arbitration scheme. Conveniently, each arbitrator arbitrates between transmission requests of the same class of service. Issued transmission requests are queued and are selectively fetched from the queues to be downstream transmitted to the network units. Preferably, one arbitrator, such as a type 1 T-CONT arbitrator, allocates data grants in a fixed manner. Another arbitrator allocates data grants in response to grouping information, to a transmission current credit, and to a class of service rules, such as bandwidth allocation rules of type 2 - 5 T-CONTs.

[00040] The invention provides a media access controller for controlling an access of a plurality of network units to a shared upstream channel, the media access controller being coupled to a receiver, for receiving grouping information from the plurality of network units. Grouping information reflects at least one parameter of fixed sized cell groups to be upstream transmitted over the shared upstream channel. The media access controller includes: (i) at least one arbitration unit, coupled between the receiver and a grant allocator, for arbitrating between requests to upstream transmit fixed sized cell groups; and (ii) a grant allocator, for selecting data grants authorizing an upstream transmission of a group of fixed sized cells in response to the arbitration.

[00041] Conveniently, the grant allocator is operative to receive allocated data grants from the at least one arbitrating unit and to select data grants in response to a predefined priority between the at least one arbitration unit. The priority is usually dependent upon the class of service. For example, the bandwidth allocation starts by allocating bandwidth (e.g.- selecting transmission requests) of type 1 T-CONTs, and ends by allocating bandwidth to type 4 T-CONTs.

[00042] The invention provides a method for allocating upstream bandwidth of a shared upstream channel of an optical network, the optical network interconnecting a headend with a plurality of network units, the method including the steps of:

[00043] Determining grouping information to be upstream transmitted from at least one network unit.

[00044] Receiving upstream transmitted grouping information.

[00045] Issuing data grants authorizing an identified network unit to transmit upstream data in response to previously received grouping information.

[00046] Conveniently, the step of issuing includes the steps of: arbitrating between requests to transmit groups of fixed sized data cells; allocating data grants in response to the arbitrating; and selecting between the allocated data grants.

[00047] Conveniently, the step of arbitrating including performing at least two arbitration cycles; and wherein the step of selecting is responsive to a predefined priorities assigned to arbitration cycles.

BRIEF DESCRIPTION OF THE DRAWINGS

- [00048] The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the drawings in which:
- [00049] Figures 1A - 1B are schematic illustration of network units, in accordance with embodiments of the invention;
- [00050] Figure 2 is a schematic illustration of a headend, in accordance with an embodiment of the invention;
- [00051] Figure 3A is a schematic illustration of an IP packet and a group of three AAL5 ATM data cells being generated from the IP packet, in accordance with an embodiment of the invention;
- [00052] Figure 3B is a schematic illustration of downstream data frames, in accordance with an embodiment of the invention;
- [00053] Figure 4 is a schematic illustration of a portion of a queues database, in accordance with an embodiment of the invention;
- [00054] Figures 5A – 5B are schematic illustrations of grouping information request vectors, in accordance with embodiments of the invention;
- [00055] Figure 5C is a schematic illustration of an upstream frame, in accordance with an embodiment of the invention; and
- [00056] Figures 6A – 6C are flow charts illustrating various arbitration mechanisms, in accordance with embodiments of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[00057] Reference is made to Figure 1A, which is a schematic illustration of a network unit NU 8, in accordance with an embodiment of the invention. Network unit can be an ONU or an OLY but this is not necessarily so.

[00058] It is noted that although the arbitration schemes depend at least partially upon grouping information, and not upon queues length, queues are still the elementary unit for exchanging information. Furthermore, a headend emulates the group of cell within each queue. Accordingly, the logical and even physical location of a queue, and especially the division between T-CONTs, ONUs and ONTs is of minor significance. For convenience of explanation it is assumed that each network unit is an ONU, that each ONU has four queues, a queue for each class of service, and that classes of service that are manageable without grouping information (such as type 1 T-CONTS) are not illustrated. Furthermore, a group of queues within an ONU, an ONT or T-CONT can be viewed by the OLT as a single queue, as long as the ONU, ONT or T-CONT has a mechanism for determining from which of its queues to fetch a group of cells in response to data grants. If is further assumed, for convenience of explanation, that a single headend, such as headend 38 of Figure 2, is coupled, via a passive optical network, to thirty two NUs such as NU 8 of Figure 1A or 1B. Media access control may not be based upon grouping information when the bandwidth is allocated regardless the presence or the parameters of data (such as in the case of type 1 T-CONT queues) or in the case of NUs that are not operable to report grouping information.

[00059] NU 8 is operable to receive two kinds of data streams via interfaces 10 and 12. It is assumed that the upstream and downstream traffic are arranged as ATM cells, although a man skilled in the art will appreciate that other communication protocols may utilized.

[00060] Interface 10 is operable to receive via input port 11 data streams that are not associated with grouping information. Interface 12 is operable to receive via input port 7 data streams that are associated with grouping information. The grouping information can be embedded within the data streams. Grouping information indicates that a group of at least one ATM cells (usually having the same Virtual Circuit and Virtual Pointer values) are to be routed as a group, preferably during consecutive upstream slots. A group of ATM cells may be generated by segmenting a single IP packet, but this is not necessarily so. The grouping information may reflect various parameters of the group of ATM cells, such as aggregate “net” ATM payload size

(e.g. - not including “idle” or stuffing signals), amount of ATM cells and/or a ratio between the “net” ATM payload and the aggregate size of the ATM cells of the group. It is noted that other parameters may be provided.

[00061] It is assumed, for convenience of explanation that the grouping information reflects the aggregate “net” ATM payload size. The aggregate size of the ATM cells of a group equal $53 * (\text{trunc}\{\text{aggregate “net” ATM payload size} / 48 \text{ byte}\} + 1)$.

[00062] It is noted that the grouping information can be compressed by implementing various compression schemes, such as lossless compression schemes, lossy compression schemes, including quantization schemes, Huffman encoding and the like.

[00063] Interface 10 provides incoming data streams, such as data streams that include ATM cells that are not AAL5 compliant (e.g.- they do not include grouping information such as IP packet length) to grouping information generator 16 that is operable to: (a) generate ATM cells, (b) generate grouping information associated with the ATM cells, and (c) determine the class of service of the ATM cells. It is noted that the determination is usually based upon the source of the data stream/ input port from which the data stream was received. Grouping information generator 16 then sends the grouping information to a queues grouping information database, determines the IP packet class and provides the ATM cells to cell distributor 18 which distributes the ATM cells among data queues DQ(1,1) – DQ(1,4) 20 – 26 according to their class of service. ATM cells of the first class of service are stored at DQ(1,1) 20, ATM cells of the second class of service are stored at DQ(1,2) 22, ATM cells of the third class of service are stored at DQ(1,3) 24, and ATM cells of the fourth class of service are stored at DQ(1,4) 26. It is noted that the amount of queues and the allocation of ATM cells among the queues can be adjusted to support various class of services and various arbitration schemes, including hierarchical arbitration schemes and the like.

[00064] If, for example, the incoming data streams are ATM data streams that include ATM cells that were generated from IP packets but are not associated with grouping information then grouping information generator 16 reassembles IP packets from the ATM cells, generates grouping information reflecting the size of the IP packet, segments the IP packets to ATM cells, provides the grouping information to queues grouping information database 30, determines the ATM cell class of service, and provides the ATM cells to cell distributor 18 which distributes the ATM cells among data queues DQ(1,1) – DQ(1,4) according to their class of service. Conveniently, the grouping information is also embedded within the group of ATM cells, such as in the case of AAL5 compliant ATM cells.

[00065] Interface 12 is operable to receive data streams that are associated with grouping information and provide them to grouping information extractor 14. Grouping information extractor 14 extracts grouping information from the received data streams, provides the grouping information to queues grouping information database 30, determines the data stream class of service and provides ATM cells to cell distributor 18 which distributes the ATM cells among data queues DQ(1,1) – DQ(1,4) according to their class of service.

[00066] At this point, queues grouping information database 30 reflects the size and location of each group of ATM cells within queues DQ(1,1) – DQ(1,4), and the queues store ATM cells of up to four class of services.

[00067] Downstream receiver 34 is operable to: (i) Receive downstream traffic being transmitted over a passive optical network. The downstream traffic includes downstream broadcast data and grants. (ii) Extract the downstream data and provides it to devices/ interfaces (not shown) that are positioned in a downstream path. (iii) Extract data grants. Data grants indicate when (during which at least one consecutive selected timeslot) to transmit upstream data from a queue out of DQ(1,1) – DQ(1,4). The extracted data grants are provided to upstream assembler and transmitter 28 that in response triggers a provision of a group of ATM cells from a queue to the passive optical network, during the at least one consecutive selected timeslot. (iv) Extract grouping information grants. Grouping information grants indicate when to transmit grouping information and what grouping information to transmit.

[00068] It is noted that grouping information grants may further include timing information for determining the timing of upstream transmission (which slot and within that slot), as a plurality of NUs may transmit their grouping information during a single timeslot. The extracted data grants are provided to upstream assembler and transmitter 28 that in response triggers a provision of a group of ATM cells from a queue to the passive optical network, during the at least one consecutive selected timeslot.

[00069] Reference is made to Figure 1B, which is a schematic illustration of a network unit NU 9, in accordance with an embodiment of the invention.

[00070] NU 9 includes a classifier 13 that has a first input 11 for receiving variable size packets such as IP packets, a marking unit 17 and a policing unit 15 located between the classifier 11 and the marking unit 17. Marking unit 17 is coupled to packet distributor 19 for providing IP packets, and is coupled to grouping information database 30 for providing grouping information. Classifier 11 analyzes the incoming IP packets to determine to which class of service they belong.

Policing unit 15 is operable to enforce policy rules. For example, the policy rules can include a mechanism for dropping packets, such as but not limited to the WRED or RED mechanisms, in response to various parameters, including network traffic load, and the like. The policy rules may also determine when to change the class of service of incoming packets. Marking unit 17 is operable to add information to the IP packets (to ‘mark’ them) in response to the classification and policing.

[00071] NU 9 further includes a reassembly unit 25 that has a second input 7 for receiving fixed sized cells, such as ATM cells. Reassembly unit 25 is coupled to grouping information extractor 27 for providing reconstructed IP packets. Grouping information extractor 27 is coupled to grouping information database 30 for providing grouping information and is coupled to packet distributor 19 for providing the reconstructed IP packets. Reassembly unit 25 has a plurality of queues (denoted VC/VP Q), each for a single combination of VP/VC fields of incoming ATM cells, such that it is able to reconstruct ATM cells.

[00072] Exemplary segmentation and reassembly operations are illustrated in Figure 3A. A segmentation process starts by receiving IP packet 100 and ends by the provision of AAL5 compliant ATM cells 112 – 116. A reassembly process starts by receiving AAL5 compliant ATM cells 112 – 116 and ends by providing IP packet 100.

[00073] An IP packet 100 is converted to an AAL5 compliant packet 110 by adding to the IP packet 100 various fields such as IP type field 102, IP packet length field 106, CRC field 104 and stuffing bytes 108. IP packet length field 106 indicates the length of IP packet 100, and can be used as a grouping information unit. The length of stuffing bytes 108 is designed such that the overall length of AAL5 compliant packet 110 equals $N * 48$ bytes, N being a positive integer. As illustrated by figure 3A, the overall length of AAL5 compliant packet 110 is $3 * 48$ bytes. AAL5 compliant packet 110 is converted to three AAL5 ATM cells 112, 114, and 116. These AAL5 ATM cells have the same VP/VC fields within their ATM headers.

[00074] Packet distributor 19 is coupled to a plurality of queues, such as PQ(1,1) – PQ(1,4) 31 – 37, and is operable to distribute the IP packets among the queues in response to the class of service of the IP packets.

[00075] PQ(1,1) – PQ(1,4) 31 – 37 are coupled to segmenting and intermediate queue unit 29, that is operable to segment IP packets to ATM cells, such as AAL5 compliant ATM cells, and to provide the ATM cells to upstream assembler and transmitter 28. It is noted that in NU 8 of figure 1A IP packets are segmented before being stored in the data queues, while in NU 9 of

figure 9 the IP packets are segmented only when they are scheduled to be upstream transmitted. It is noted that the order of the storing and segmenting operations is not significant.

[00076] Segmenting and intermediate queue unit 29 is operable to segment the IP packets to ATM cells, and to store them in intermediate queues. In cases where not all the ATM cells of an ATM cell group that originated from a single IP packet can be transmitted, the remaining ATM cells, can be provided to upstream assembler and transmitter 28 during another timeslot.

[00077] Referring to Figure 2, illustrating headend, such as OLT 38, in accordance with an embodiment of the invention.

[00078] OLT 38 includes (i) downstream data interface 40, (ii) grant controller 48, (iii) downstream transmitter 46, (iv) grant allocator 49, (v) upstream receiver 54, and (vi) grant queues GQ1 – GQ4 40 – 46.

[00079] Downstream data interface 40 and grant allocator 49 are coupled to downstream transmitter 48. Grant controller 48 is coupled to upstream receiver 54 and to grant queues GQ1 – GQ4 40 – 46. Grant queues GQ1 – GQ4 40 – 46 are coupled to grant allocator 49.

[00080] Downstream data interface is operable to receive data to be downstream transmitted to NUs over a passive optical network and to provide said data to downstream transmitter 46. Downstream transmitter 46 is operable to further receive data grants and grouping information request grants from grant allocator 48 and to generate downstream frames.

[00081] Upstream receiver 54 receives upstream transmitted frames and is operable to extract grouping information out of the upstream transmitted frames, to provide the grouping information to grant allocator 48 and to provide the upstream data along an upstream path from OLT 38 to other headends, networks or other network units.

[00082] Grant controller 48 is operable to receive grouping information and to issue data grants in response to the grouping information. The allocated data grants are stored in grant queues GQ1 – GQ4 40 – 46, according to their class of service.

[00083] Grant queues GQ1 – GQ4 40 – 46 are coupled to grant allocator 49 which selectively fetches issued data grants according to a predefined order, and sends a sequence of data grants to downstream transmitted 46, where the data grants are assembled in to a downstream frame. For example, the selection may start by selecting a issued data grants of the highest class of service, and continue to the lower priority class of service grant queues, as long as additional timeslots may be allocated. It is noted that when implementing a fixed timeslot allocation, such as in the case of

type 1 T-CONTs, an additional unit (not shown) may be operable to generate data grants to allows a fixed timeslot allocation.

[00084] Reference is now made to Figure 4 which illustrates an exemplary portion of grant controller 48. Grant controller 48 includes grouping information database 50, a plurality of arbitration units such as arbitrators 61 – 64, and data converter 65.

[00085] Portion 51_1 of grouping information database 50 stores the grouping information of DQ(1,1) - a list of “net” sizes/length of relevant payload of seventy consecutive ATM cell groups that are stored at DQ(1,1): SZ(1,1,1) – SZ(1,1,70). SZ(1,1,1) being the size of the group of ATM cells that is located at the head of DQ(1,1). Portion 51_32 of grouping information database 50 stores the grouping information of DQ(32,1) - a list of “net” sizes of sixty consecutive ATM cell groups that are stored at DQ(32,1): SZ(32,1,1) – SZ(32,1,60). Portions 51_1 – 51_32 of queues database 50 that are associated with a first class of service are accessible to a first arbitrator 61 that determines which first class of service ATM cell group to transmit.

[00086] Portion 52_1 of grouping information database 50 stores the grouping information of DQ(1,2) - a list of “net” sizes of fifty four consecutive ATM cell groups that are stored at DQ(1,2): SZ(1,2,1) – SZ(1,2,54). SZ(1,2,1) being the size of the group of ATM cells that is located at the head of DQ(1,2). Portion 52_32 stores the grouping information of DQ(32,2) - a list of “net” sizes of sixty three consecutive ATM cell groups that are stored at DQ(32,2): SZ(32,2,1) – SZ(32,2,63). Portions 52_1 – 52_32 of queues database 50 are associated with a second class of service are accessible to a second arbitrator 62 that determines which second class of service ATM cell group to transmit.

[00087] Portion 53_1 of grouping information database 50 stores the grouping information of DQ(1,3) - a list of “net” sizes of twenty five consecutive ATM cell groups that are stored at DQ(1,3): SZ(1,3,1) – SZ(1,3,25). SZ(1,3,1) being the size of the group of ATM cells that is located at the head of DQ(1,3). Portion 53_32 stores the grouping information of DQ(32,3) - a list of “net” sizes of fifteen consecutive ATM cell groups that are stored at DQ(32,3): SZ(32,3,1) – SZ(32,3,15). Portions 52_1 – 52_32 of queues database 50 are associated with a third class of service are accessible to a third arbitrator 63 that determines which third class of service ATM cell group to transmit.

[00088] Portion 54_1 of grouping information database 50 stores the grouping information of DQ(1,4) - a list of “net” sizes of sixty consecutive ATM cell groups that are stored at DQ(1,4): SZ(1,4,1) – SZ(1,4,60). SZ(1,4,1) being the size of the group of ATM cells that is located at the

head of DQ(1,4). Portion 54_32 stores the grouping information of DQ(32,4) - a list of “net” sizes of ninety consecutive ATM cell groups that are stored at DQ(32,4): SZ(32,4,1) – SZ(32,4,90). Portions 54_1 – 54_32 of queues database 50 are associated with a fourth class of service are accessible to a fourth arbitrator 64 that determines which fourth class of service ATM cell group to transmit.

[00089] The results of the arbitration cycles are fed to data converter 65 and then to grant queues GQ1 – GQ4 40 - 46. Data converter 65 is operable to convert a result of an arbitration, that usually reflects the “net” size of ATM payload to transmit, to an amount of timeslot to allocate to the transmission. Accordingly, GQ1 - GQ4 40 – 46 store the amount of consecutive timeslots to allocate for an upstream transmission from a certain NU queue. These amounts are also referred to as issued data grants.

[00090] According to an aspect of the invention, OLT 38 determines the amount of grouping information to receive from each NU. This determination allows for reducing the amount of upstream bandwidth assigned to the transmission of grouping information.

[00091] It is noted that during a MAC cycle a maximal predefined amount of data may be upstream transmitted from each queue. The maximal amount is limited by upstream transmission rates and by the rate in which the upstream data is provided to the queue. The maximal amount can be predefined or can be dynamically changed to reflect various changes in the network.

[00092] At each MAC cycle OLT 38 determines the grouping information to be upstream transmitted by the NUs. The amount of grouping information is responsive to the grouping information that has already been received by OLT 38, the maximal predetermined amount and an estimation of the ATM cell groups that are stored at the queues of which their grouping information was not yet transmitted to OLT 38.

[00093] An estimation of ATM cell groups size that are supposed to be stored in a data queue can be based upon statistical analysis of ATM cell groups received from (i) that data queue, (ii) queues that belong to the same T-CONT, (iii) queues that belong to the same ONU, (iv) queues that service the same class of service, from ATM cell groups received from a certain application, from ATM cell groups that originated from packets of a certain networking protocol, from ATM groups originating from the same source and/or destined to the same destination, and the like.

[00094] For example, and referring to Figure 4, it is assumed that (a) the grouping information is the length (in bytes) of IP packets from which ATM cells were generated (b) the received data streams are Ethernet compliant and the minimal size of an Ethernet packet is 64

bytes, (c) the maximal amount of bytes that can be transmitted from data queue DQ(32,4) during a single MAC cycle is 5000, (d) portion 54_32 stores SZ(32,4,1) – SZ(32,4,90), (e) the aggregate size of SZ(32,4) – SZ(32,4,90) is 4540. Then OLT 38 will request ONU to report the sizes of the next ten ATM cell groups within DQ(32,4) : (5000–4540)/64 = 10.

[00095] Referring to figure 3A illustrating downstream frames 120 - 140, in accordance to embodiments of the invention.

[00096] The upper part of Figure 3A illustrates three downstream frames, such as T-frames, that are transmitted during a single MAC cycle. Each T-frame out of T-frames 120, 130 and 140 includes fifty-three slots.

[00097] T-frame 120 includes fifty-one ATM cell slots and two framing slots, also referred to as PLOAM slots. A first PLOAM slot 121 is followed by twenty seven ATM slots 122_2 – 122_28, second PLOAM slot 123 and twenty four ATM slots 122_30 – 122_53.

[00098] T-frame 120 is followed by T-frame 130 that includes fifty-one ATM cell slots and two framing slots, also referred to as PLOAM slots. A first PLOAM slot 131 is followed by twenty seven ATM slots 132_2 – 132_28, second PLOAM slot 133 and twenty four ATM slots 132_30 – 132_53.

[00099] T-frame 130 is followed by T-frame 140 that includes forty-three ATM cell slots, two framing (PLOAM) slots and eight GI slots. A first PLOAM slot 141 is followed by twenty seven ATM slots 142_2 – 142_28, second PLOAM slot 143, sixteen ATM slots 142_30 – 142_45, and eight GI slots 144_1 – 144_8 in which a grouping information request vector 146 is transmitted.

[00100] The mentioned above format is suited to an optical passive network in which the downstream data rate equals the data upstream rate. Accordingly, each pair of PLOAM slots (121, 123), (131, 133) and (141,143) determines which upstream data to transmit during an upstream frame of fifty three slots.

[00101] Eight GI slots 144_1 – 144_8 are allocated for transmitting grouping information requests that determine which grouping information to transmit upstream during a MAC cycle of three T-frames.

[00102] The lower part of Figure 3B illustrates a MAC cycle the includes three T-frames, that is suited to an optical passive network in which the downstream data rate is three times faster than the upstream data rate. In such a case portions 146_1, 146_2, 146_3 and 146_4 of grouping

information request vector 146 are transmitted during slots 131, 133, 141 and 143 accordingly. It is noted that the third T-frame (140) can be utilized to convey more upstream ATM cells or can include more grouping information request slots.

[000103] Referring to Figure 5A illustrating grouping information request vectors 152, 153 and 154 in accordance to embodiments of the invention.

[000104] Grouping information request vector 153 includes two types of fields - a first type field that indicates an identity of a queue and a second type indicating the amount of grouping information, relating to that queue, to be transmit. For example, field 153_1 “REQUEST TO REPORT GI RELATING TO THE Y1 QUEUE” is followed by field 153_2 “REQUEST TO REPORT GI OF X1 GROUPS”. These two fields are interpreted as a request to upstream transmit grouping information relating to X1 consecutive ATM cell groups that are stored at the Y1'th queue, starting from the first ATM cell group which its grouping information was not reported yet. Each NU monitors the grouping information of each of the NU queues and tracks the previously upstream transmitted grouping information. Accordingly, each NU is able to determine which grouping information to prepare for upstream transmission.

[000105] Grouping information request vector 152 includes only the second type of fields and may be transmitted when the order of the relevant queues is predefined.

[000106] Grouping information request vector 154 is compressed according to a predefined compression scheme. It includes compressed fields of the second type only, although this is not necessarily so. For example, a compressed grouping information request vector may also include the two types of fields.

[000107] Grouping information request vector 153 includes a request to transmit grouping information relating to X1 ATM cell groups from queue Y1 (fields 153_1 and 153_2), a request to transmit grouping information relating to X2 ATM cell groups from queue Y2 (fields 153_3 and 153_4), and a request to transmit grouping information relating to X3 ATM cell groups from queue Y3 (fields 153_5 and 153_6).

[000108] Grouping information request vector 152 includes a request to transmit grouping information relating to X1 ATM cell groups from queue Y1 (field 152_1, the identity of the queue is predefined), a request to transmit grouping information relating to X2 ATM cell groups from queue Y2 (field 152_2, the identity of the queue is predefined), a request to transmit grouping information relating to X3 ATM cell groups from queue Y3 (fields 152_3, the identity

of the queue is predefined) and a request to transmit grouping information relating to X4 ATM cell groups from queue Y4 (fields 152_4 , the identity of the queue is predefined).

[000109] Grouping information request vector 154 includes the same requests as grouping information request vector 152, but in a compressed form.

[000110] Reference is now made to Figure 5B, illustrating grouping information vectors 160 and 170 in accordance to embodiments of the invention.

[000111] Grouping information vector 170 is a compressed version of grouping information vector 160.

[000112] The timing of the upstream transmission of each portion of grouping information vectors is determined by OLT 38 and may either be predefined to included within grouping information request vectors. It is noted the methods for timing the upstream transmission may implemented in a manner that resembles the timing of ITU-T Recommendation G.983.4 compliant mini-slots.

[000113] Grouping information vectors 160 and 170 are transmitted in response to (i) a reception of fields 152_1 and 152_2 of grouping information request vector 152, or (ii) a reception of fields 154_1 and 154_2 of grouping information request vector 154, or (iii) a reception of fields 153_1, 153_2, 153_3 and 153_4 of grouping information request vector 153. Accordingly, the grouping information vector 160 starts by X1 fields 162_1 - 162_X1 that embeds the grouping information relating to X1 ATM cell groups being stored within queue Y1. These fields are followed by X2 fields 163_1 – 163_X2 that embed the grouping information relating to X2 ATM cell groups being stored within queue Y2.

[000114] Grouping information vector 170 includes X1 fields 172_1 – 172_X1 that embed compressed grouping information relating to X1 ATM cell groups being stored within queue Y1. These fields are followed by X2 fields 173_1 – 173_X2 that embed compressed grouping information relating to X2 ATM cell groups being stored within queue Y2.

[000115] It is noted that the timing of the upstream transmission may be predetermined, negotiated between the OLT 38 and NUs or be included within the grouping information request vector.

[000116] Figure 5C illustrates an upstream frame, in accordance with an embodiment of the invention.

[000117] The upstream frame includes three upstream, such as T-frames 180, 190 and 200.

Each T-frame includes 53 slots.

[000118] Portions of grouping information vectors (also referred to as GIV), are transmitting in predefined slots, such as GIVs 210_1, 210_2 and 210_3 OLT 38 determines the position (timing) of each of these slots and also determines which grouping information to send during each slot.

Each T-frame out of T-frames 180, 190 and 200 starts by a GI vector slot, but this is not necessarily so. The inclusion of a GI vector slot within each T-frame allows to update the GI of queues several times during a MAC cycle, thus allowing relatively frequent update of the timing information, and especially allows for upstream transmitting grouping information of high priority (high class of service), thus allowing to issue upstream transmit the high priority information before the current MAC cycle ends.

[000119] The length of each GIV slot may be predefined but it also may be configured to match the length of grouping information vectors from a certain set of queues. For example, the length (in ATM cells) of 210_1, 210_2 and 210_3 are (Z1-1), (Z2-1) and (Z3-1) accordingly.

[000120] Reference is now made to Figures 6A – 6C, illustrating various arbitration mechanisms, in accordance with preferred embodiments of the invention. These arbitration schemes may be implemented by at least one of arbitrator out of arbitrators 61 – 64.

[000121] Conveniently, each arbitrator is operable to allocate resources/ grant upstream data transmission of data of a certain service of class. Arbitrator 61 is operable to grant upstream data transmissions of the first class of service, arbitrator 62 is operable to grant upstream data transmissions of the second class of service, arbitrator 63 is operable to grant upstream data transmissions of the third class of service, arbitrator 64 is operable to grant upstream data transmissions of the fourth class of service.

[000122] It is noted that the result of an arbitration process, such as an arbitration process implemented by arbitrators 61 – 64, is a data grant. The data grants are further selectively fetched by grant allocator 49, that actually allocates upstream timeslots.

[000123] Figure 6A is a flow chart illustrating a first arbitration mechanism 310. First arbitration mechanism 310 is operable to allocate “fixed” data grants - data grants for transmitting a fixed amount of data. First arbitration mechanism 310 starts by a step 311 of defining an the fixed amount of slot quota to allocate to each queue out of a set of queues that are arbitrated by first arbitration process 310. Usually, the set include queues of the same class of service. It is noted that queues that are associated to the same class of service may be arbitrated

by a single arbitration process, but this is not necessarily so and they may be partitioned to a plurality of sets, each set of a predefined priority. It is further noted that each queue may be assigned a different fixed amount of slot quota.

[000124] Step 311 is followed by step 312 of allocating data grants in accordance with definitions of step 311. The data grant allocation is usually done regardless the emptiness of the queues. Step 312 may be executed each MAC cycle.

[000125] Figure 6B is a flow chart illustrating a second arbitration mechanism 320. Second arbitration mechanism 320 is operable to allocate a predefined amount of data grants in response to an indication that a queue that takes part in the arbitration cycle is not empty. Such a queue is referred to as a participating queue. Second arbitration mechanism 320 starts by a step 321 of defining an amount of slot quota to allocate to each participating queue during an arbitration cycle. Step 321 is followed by step 322 of receiving an indication of the participating queues (which queues out of possible participating queues are not empty). This indication may be derived from previously received grouping information. Step 322 is followed by a step 324 of allocating data grants to participating queues.

[000126] Figure 6C is a flow chart illustrating a third arbitration mechanism 330. Third arbitration mechanism 330 is operable to allocate data grants in response to a predefined credit assigned to each queue and to the grouping information.

[000127] Third arbitration mechanism 330 starts by a step 332 of defining a “credit” quote for each queue that is arbitrated by third arbitration mechanism 330. The “credit” quote is an amount of timeslots/data grants that can be allocated to upstream data transmission from the queue. The credit quote assigned to each queue reflects the queue priority/ relative allocation of data grants. According to an aspect of the invention that the credit quota and/or can be dynamically changed. Step 332 is followed by step 334 of receiving grouping information relating to queues that are arbitrated by third arbitration mechanism 330.

[000128] Step 334 is followed by step 336 of determining from which queue to start the current arbitration cycle. Conveniently, the queues are arranged in a cyclical manner and an arbitration cycle starts from a queue that follows the last queue that was processed during a previous arbitration cycle.

[000129] Step 336 is followed by step 338 of determining whether the current queue is empty. If the answer is “yes” step 338 is followed by step 340 of selecting a new current queue. If the

queues are arranged in a round robin formation, the next current queue is an adjacent queue of the round robin formation. If the queue is not empty, step 338 is followed by step 342.

[000130] Step 342 includes determining the amount of ATM cell groups that can be transmitted given the current credit of the current queue. It is noted that during a first iteration of the arbitration mechanism the current credit may equal the credit quote, but as further illustrated, the current quote is updated during the arbitration process. The ATM cell group may include consecutive ATM cell groups, starting with the first ATM cell group – the ATM cell group that is stored at the head of the current queue. According to an aspect of the invention step 342 includes determining a “net” amount of data that may be transmitted, whereas the “net” amount is further converted to an amount of data cells.

[000131] If the current credit does not allow for transmitting at least the first ATM cell group then step 342 is followed by step 341 of increasing the current credit by the credit quota. Step 341 is followed by step 340.

[000132] If the current credit does allow for transmitting a set of consecutive ATM groups, step 342 is followed by step 344 of allocating data grants for the transmission of the set and updating the current credit by adding a credit quota and decreasing the aggregate size of the ATM cell groups that form the set.

[000133] Step 344 is followed by step 346 of determining when the arbitration cycle ends. The determination can be responsive to the amount of data grants that was allocated or to the amount of iterations of steps 338-342. The overall amount of allowed data grants can be dynamically changed.

[000134] If the arbitration cycle ended, step 346 is followed by step 350 ending the arbitration cycle. Else, step 346 is followed by step 340 of selecting a new current queue.

[000135] Distinct arbitration schemes, such as but not limited to arbitration schemes 310 – 330, may be implemented by various arbitrators, thus allowing for tailoring arbitration schemes to the characteristics of each class of service. These characteristics may include, bursty behavior patterns, delay sensitivity and the like.

[000136] Systems and methods that are adapted to handle various class of services include dependent arbitrators and independent arbitrators. Independent arbitrators allocate data grants regardless the allocation of data grants by other arbitrators. Dependent arbitrators allocate data grants in response to data grant allocated by other arbitrators. Independent arbitrators are usually utilized for the highest (most important) class of services and to fixed arbitration schemes.

Dependent arbitrators are usually utilized for lower class of services. Dependent arbitrators may receive various input from other arbitrators, such as a residual overall data grant amount to be allocated during an arbitration cycle, or a relationship between allocated data grants to queues. In dependent arbitrators that implement the third arbitration mechanism 330 the credit quota, the current credit and the total amount of data grant location per arbitration cycle can be dynamically changed in response to the receive input from other arbitrators.

[000137] For example, systems and methods that are operable to handle the following class of services (provisioned bandwidth, guaranteed bandwidth, assured bandwidth and non-assured bandwidth) may (i) utilize an independent arbitrator that implements a first arbitration mechanism 310 for handling provisioned bandwidth, (ii) utilize a dependent arbitrator that depends upon the implements a second arbitration mechanism 320 for handling allocated bandwidth, the dependent arbitrators receives the amount of overall allowed data grants, given the results of the independent arbitrator, (iii) utilize a dependent arbitrator that implements a third arbitration mechanism 330 for handling provisioned bandwidth, the dependent arbitrators receives the amount of overall allowed data grants, given the results of the independent arbitrator and the mentioned above dependent arbitrator. The dependent arbitrator can also receive a relationship between data grant allocations of distinct queues in response to the allocation of the previous arbitrator, and (iv) utilize a dependent arbitrator that implements a third arbitration mechanism 330 for handling non-assured bandwidth.

[000138] It is noted that other combinations of arbitration schemes may be implemented to fit various class of services, such as but not limited to provisioned bandwidth, minimum latency, assured bandwidth, non-assured bandwidth, minimum drop, minimum jitter and a combination of the mentioned above class of services.

[000139] An example of the amount of grouping information slots and other system parameters, is illustrated below.

[000140] It is assumed that: (i) a MAC cycle includes three T-frame, (ii) downstream data rate of 622 Mbytes/sec, (iii) upstream data rate of 155 Mbytes/sec, (iv) an IP packet is represented by two to five ATM cells; (v) a network unit that is connected to 100baseT communication lines receives a maximal amount of thirty IP packets, each being 64 byte long, per MAC cycle, it typically receives ten IP packets, each being 200 byte long,(vi) a network unit that is connected to 10baseT communication lines receives a maximal amount of three IP packets, each being 64 byte long, per MAC cycle, it typically receives a single 200 bytes long IP packet, (vii) a plurality of network units share the same upstream channel, about four to eight of them are connected to

100baseT communication lines, the other network units are connected to 10baseT communication lines.

[000141] Accordingly, during a single MAC cycle a maximal amount of 318 IP packets are received, typically only 129 IP packets are received. A grouping information unit is a byte long, but may also be two bytes long. Assuming that the transmission of grouping information requires additional bytes, for example, for re-transmitting grouping information, for transmitting CRC bits and the like, then each MAC cycle an amount of $3*318*1$ bytes, that can be encapsulated within 20 ATM cells, are required.

[000142] It is noted that the method and apparatus according to the present invention can be implemented either in hardware, in software or in a combination thereof.

[000143] It will be apparent to those skilled in the art that the disclosed subject matter may be modified in numerous ways and may assume many embodiments other than the preferred form specifically set out and described above.

[000144] Accordingly, the above disclosed subject matter is to be considered illustrative and not restrictive, and to the maximum extent allowed by law, it is intended by the appended claims to cover all such modifications and other embodiments, which fall within the true spirit and scope of the present invention.

[000145] The scope of the invention is to be determined by the broadest permissible interpretation of the following claims and their equivalents rather than the foregoing detailed description.